Available online at www.sciencedirect.com
SCIENCE DIRECT®
ELSEVIER
Automatica 40 (2004) 639–646
automatica
www.elsevier.com/locate/automatica

Brief paper

Unconstrained optimal control of regular languages☆

Jinbo Fu, Asok Ray*, Constantino M. Lagoa

Mechanical Engineering Department, The Pennsylvania State University, 137 Reber Building, University Park, PA 16802-1412, USA

Received 14 April 2002; received in revised form 27 March 2003; accepted 3 November 2003

Abstract

This paper formulates an unconstrained optimal policy for control of regular languages realized as deterministic finite state automata (DFSA). A signed real measure quantifies the behavior of controlled sublanguages based on a state transition cost matrix and a characteristic vector as reported in an earlier publication. The state-based optimal control policy is obtained by selectively disabling controllable events to maximize the measure of the controlled plant language without any further constraints. Synthesis of the optimal control policy requires at most n iterations, where n is the number of states of the DFSA model. Each iteration solves a set of n simultaneous linear algebraic equations. As such, computational complexity of the control synthesis is polynomial in n.

Keywords: Discrete event systems; Optimal control; Finite state automata; Language measure

1. Introduction

The discrete-event dynamic behavior of physical plants is often modeled as regular languages that can be realized by finite-state automata (Ramadge & Wonham, 1987). The sublanguage of a controlled plant could be different under different supervisors that are constrained to satisfy different specifications. Such a partially ordered set of sublanguages requires a quantitative measure for total ordering of their respective performance. To address this issue, Wang and Ray (2002) have developed a signed measure of regular languages. This work was followed by Surana and Ray (2003) who have constructed a vector space of sublanguages with a metric based on the total variation measure of the language.

Several researchers have proposed optimal control of deterministic finite state automata (DFSA) based on different assumptions. Some of these researchers have attempted to quantify the controller performance using different types of cost assigned to the individual events. Pasino and Antsaklis (1989) proposed path costs associated with state transitions and hence optimal control of a discrete event system is equivalent to following the shortest path on the graph representing the uncontrolled system. Kumar and Garg (1995) introduced the concept of payoff and control costs that are incurred only once regardless of the number of times the system visits the state associated with the cost. Consequently, the resulting cost is not a function of the dynamic behavior of the plant. Brave and Heyman (1990) introduced the concept of optimal attractors in discrete-event control. Sengupta and Lafortune (1998) used control cost in addition to the path cost in optimization of the performance index for trade-off between finding the shortest path and reducing the control cost. Although costs were assigned to the events, no distinction was made for events generated at (or leading to) different states that could be “good” or “bad”. These optimal control strategies have addressed performance enhancement of discrete-event control systems without a quantitative measure of languages.

This paper introduces the concept of unconstrained optimal control of DFSA based on a specified language measure. Starting with the (regular) language of the unsupervised (i.e., open loop) plant, the optimal control policy maximizes the performance of a sublanguage without any further constraints. The paper is organized in six sections.

☆ This paper was not presented at any IFAC Meeting. This paper was recommended for publication in revised form by Associate Editor Xiren Cao under the direction of Editor Tamer Başar. This work has been supported in part by the Army Research Office under Grant No. DAAD19-01-1-0646; National Science Foundation under Grant No. ECS-9984260; and NASA Glenn Research Center under Grant No. NAG3-2448 and NAG3-2940.

* Corresponding author. Tel.: +1-814-865-6377; fax: +1-814-863-4848.

E-mail addresses: jinbofu@adelphia.net (J. Fu), axr2@psu.edu (A. Ray), lagoa@engr.psu.edu (C.M. Lagoa).

0005-1098/$ - see front matter © 2003 Elsevier Ltd. All rights reserved.
doi:10.1016/j.automatica.2003.11.011

Section 2 briefly describes the language measure and introduces the notation used in the sequel. Section 3 presents the underlying theory of the optimal control policy in the context of the language measure. Section 4 constructs a procedure for synthesis of discrete-event optimal controllers and shows that computational complexity of the control synthesis procedure is polynomial in the number of the DFSA states. Section 5 presents an application example for discrete-event optimal control of a gas turbine engine. The paper is summarized and concluded in Section 6 along with recommendations for future research.

2. Brief review of the language measure

This section briefly reviews the concept of signed real measure of regular languages (Wang & Ray, 2002). Let the plant behavior be modeled as a DFSA $G_i \equiv (Q, \Sigma, \delta, q_i, Q_m)$, where Q is the finite set of states with $|Q| = n$ excluding the dump state (Ramadge & Wonham, 1987) if any, and $q_i \in Q$ is the initial state; Σ is the (finite) alphabet of events; Σ^* is the set of all finite-length strings of events including the empty string ε; the (possibly partial) function $\delta : Q \times \Sigma \to Q$ represents state transitions and $\delta^* : Q \times \Sigma^* \to Q$ is an extension of δ; and $Q_m \subseteq Q$ is the set of marked states.

Definition 1. A DFSA G_i, initialized at $q_i \in Q$, generates the language $L(G_i) \equiv \{s \in \Sigma^* : \delta^*(q_i, s) \in Q\}$ and its marked sublanguage $L_m(G_i) \equiv \{s \in \Sigma^* : \delta^*(q_i, s) \in Q_m\}$.

The language $L(G_i)$ is partitioned as the non-marked and the marked languages, $L^o(G_i) \equiv L(G_i) - L_m(G_i)$ and $L_m(G_i)$, consisting of event strings that, starting from $q_i \in Q$, terminate at one of the non-marked states in $Q - Q_m$ and one of the marked states in Q_m, respectively. The set Q_m is partitioned into Q_m^+ and Q_m^-, where Q_m^- contains all *good* marked states that we desire to reach and Q_m^- contains all *bad* marked states that we want to avoid, although it may not always be possible to avoid the bad states while attempting to reach the good states. The marked language $L_m(G_i)$ is further partitioned into $L_m^+(G)$ and $L_m^-(G_i)$ consisting of good and bad strings that, starting from q_i, terminate on Q_m^+ and Q_m^-, respectively.

A signed real measure $\mu : 2^{\Sigma^*} \to \Re \equiv (-\infty, \infty)$ is constructed for quantitative evaluation of every event string $s \in \Sigma^*$. The language $L(G_i)$ is decomposed into null (i.e., $L^o(G_i)$), positive (i.e., $L_m^+(G_i)$), and negative (i.e., $L_m^-(G_i)$) sublanguages.

Definition 2. The language of all strings that, starting at a state $q_i \in Q$, terminates on a state $q_j \in Q$, is denoted as $L(q_i, q_j)$. That is, $L(q_i, q_j) \equiv \{s \in L(G_i) : \delta^*(q_i, s) = q_j\}$.

Definition 3. The characteristic function that assigns a signed real weight to state-partitioned sublanguages $L(q_i, q_j)$ is defined as: $\chi : Q \to [-1, 1]$ such that

$$\chi(q_j) \in \begin{cases} [-1, 0) & \text{if } q_j \in Q_m^- \\ \{0\} & \text{if } q_j \notin Q_m \\ (0, 1] & \text{if } q_j \in Q_m^+ \end{cases}.$$

Definition 4. The event cost is conditioned on a DFSA state at which the event is generated, and is defined as $\tilde{\pi} : \Sigma^* \times Q \to [0, 1]$ such that $\forall q_j \in Q$, $\forall \sigma_k \in \Sigma$, $\forall s \in \Sigma^*$,

- $\tilde{\pi}(\sigma_k \mid q_j) = 0$ if $\delta(q_j, \sigma_k)$ is undefined; $\tilde{\pi}[\varepsilon \mid q_j] = 1$;
- $\tilde{\pi}(\sigma_k \mid q_j) \equiv \tilde{\pi}_{jk} \in [0, 1)$; $\quad \sum_k \tilde{\pi}_{jk} < 1$;
- $\tilde{\pi}(\sigma_k s \mid q_j) = \tilde{\pi}(\sigma_k \mid q_j) \tilde{\pi}(s \mid \delta(q_j, \sigma_k))$.

Now we define the measure of any sublanguage of the $L(G_i)$ in terms of the signed characteristic function χ and the non-negative event cost $\tilde{\pi}$.

Definition 5. The signed real measure μ of a singleton string set $\{s\} \subset L(q_i, q_j) \subseteq L(G_i) \in 2^{\Sigma^*}$ is defined as

$$\mu(\{s\}) \equiv \chi(q_j) \tilde{\pi}(s \mid q_i) \quad \forall s \in L(q_i, q_j).$$

The signed real measure of $L(q_i, q_j)$ is defined as

$$\mu(L(q_i, q_j)) \equiv \sum_{s \in L(q_i, q_j)} \mu(\{s\})$$

and the signed real measure of a DFSA G_i, initialized at the state $q_i \in Q$, is denoted as

$$\mu_i \equiv \mu(L(G_i)) = \sum_j \mu(L(q_i, q_j)).$$

Definition 6. The state transition cost of the DFSA is defined as a function $\pi : Q \times Q \to [0, 1)$ such that $\forall q_j, q_k \in Q$, $\pi(q_k \mid q_j) = \sum_{\sigma \in \Sigma : \delta(q_j, \sigma) = q_k} \tilde{\pi}(\sigma \mid q_j) \equiv \pi_{jk}$ and $\pi_{jk} = 0$ if $\{\sigma \in \Sigma : \delta(q_j, \sigma) = q_k\} = \emptyset$. The $n \times n$ state transition cost matrix, denoted as Π-matrix, is defined as

$$\Pi = \begin{bmatrix} \pi_{11} & \pi_{12} & \cdots & \pi_{1n} \\ \pi_{21} & \pi_{22} & \cdots & \pi_{2n} \\ \vdots & & \ddots & \vdots \\ \pi_{n1} & \pi_{n2} & \cdots & \pi_{nn} \end{bmatrix}.$$

Proposition 1. *Given a state transition cost matrix $\Pi \in \Re^{n \times n}$, the operator $[I - \Pi]$ is invertible and the bounded linear operator $[I - \Pi]^{-1} \geqslant 0$ where the matrix inequality is implied elementwise.*

Proof. It follows from Definitions 4 and 6 that $\exists\, \theta \in (0, 1)$ such that the induced infinity norm $\|\Pi\|_\infty \equiv \max_i \Sigma_j \pi_{ij} = 1 - \theta$. Then, $[I - \Pi]^{-1}$ is invertible and is a bounded linear operator with $\|[I - \Pi]^{-1}\|_\infty \leqslant \theta^{-1}$ (Naylor & Sell, 1982, p. 431). Using Taylor series expansion, $[I - \Pi]^{-1} = \sum_{k=0}^{\infty} [\Pi]^k$. Since each element of Π is non-negative, so is each element of $[\Pi]^k$. Hence, $[I - \Pi]^{-1} \geqslant 0$ elementwise. □

Wang and Ray (2002) have shown that the measure $\mu_i \equiv \mu(L(G_i))$ of the language $L(G_i)$ can be expressed as $\mu_i = \sum_j \pi_{ij}\mu_j + \chi_i$ where $\chi_i \equiv \chi(q_i)$. Equivalently, in vector notation: $\mu = \Pi\mu + X$, where the measure vector $\mu \equiv [\mu_1\mu_2\cdots\mu_n]^T$ and the characteristic vector $X \equiv [\chi_1\chi_2\cdots\chi_n]^T$. By Proposition 1, the measure vector μ is uniquely determined as: $\mu = [I - \Pi]^{-1}X$.

3. Derivation of the optimal control policy

This section presents the theoretical foundations of the unconstrained optimal control of discrete event systems. Let $\mathscr{S} \equiv \{S^0, S^1, \ldots, S^N\}$ be the finite set of all supervisory control policies for the unsupervised plant automaton G, where S^0 is the null controller, i.e., no event is disabled and hence $L(S^0/G) = L(G)$. Therefore, the controller cost matrix $\Pi(S^0) = \Pi^0 \equiv \Pi^{\text{plant}}$ that is the Π-matrix of the unsupervised plant automaton G. For a supervisor $S^i, i \in \{1, 2, \ldots, N\}$, the control policy selectively disables certain controllable events by which the corresponding elements of the $\tilde{\Pi}$-matrix become zero. The (elementwise) inequality holds: $\Pi(S^k) \equiv \Pi^k \leqslant \Pi^0$ where $L(S^k/G) \subseteq L(G)\ \forall S^k \in \mathscr{S}$. The ijth element of Π^k is denoted as π_{ij}. The performance vector is denoted as $\mu(S^k) \equiv \mu^k = [I - \Pi^k]^{-1}X$ and its jth element as μ_j^k.

Definition 7. For any supervisor $S \in \mathscr{S}$ and any measure vector $v \in \mathfrak{R}^n$, the affine operator $T(S): \mathfrak{R}^n \to \mathfrak{R}^n$ is defined as: $T(S)v = \Pi(S)v + X$.

Proposition 2. *$\forall S \in \mathscr{S}, T(S)$ is a contraction operator, and there exists a unique measure vector $\mu(S)$ such that $\mu(S) = T(S)\mu(S)$.*

Proof. Since Π^0 is a contraction operator, $0 \leqslant \Pi(S) \leqslant \Pi^0$, and X is a constant vector, $T(S)$ is also a contraction operator. Hence, $\exists$ a unique fixed point $\mu(S) = T(S)\mu(S)$ in the Banach space $\mathfrak{R}^n$ (Naylor & Sell, 1982, p. 126). □

Corollary 1 to Proposition 2. The unique fixed point of the contraction operator $T^k \equiv T(S^k)$ is: $\mu^k = [I - \Pi^k]^{-1}X$.

Proof. The unique fixed point μ^k of T^k satisfies the identity $\mu^k = \Pi^k\mu^k + X$. As $0 \leqslant \Pi^k \leqslant \Pi^0$ elementwise, we have $\|\Pi^k\|_\infty \leqslant \|\Pi^0\|_\infty < 1$. Hence, the operator $[I - \Pi^k]^{-1}$ is bounded. □

Corollary 2 to Proposition 2. The operator $[I - \Pi^k]$ has a real positive determinant, i.e., $Det[I - \Pi^k] > 0$.

Proof. Eigenvalues of the real matrix Π^k are located within the unit circle and they appear as real or complex conjugates. Therefore, eigenvalues of $[I - \Pi^k]$ have positive real parts. So, $Det[I - \Pi^k]$ is real positive. □

Proposition 3. (i) *Let j be such that $\mu_j^k = \min_{\ell \in \{1,2,\ldots,n\}} \mu_\ell^k$; if $\mu_j^k \leqslant 0$, then $\chi_j \leqslant 0$, and if $\mu_j^k < 0$, then $\chi_j < 0$;* (ii) *Let j be such that $\mu_j^k = \max_{\ell \in \{1,2,\ldots,n\}} \mu_\ell^k$; if $\mu_j^k \geqslant 0$, then $\chi_j \geqslant 0$, and if $\mu_j^k > 0$, then $\chi_j > 0$.*

Proof. (i) The DFSA satisfies the identity $\mu_j^k = \sum_{\ell \in \{1,2,\ldots,n\}} \pi_{j\ell}^k \mu_\ell^k + \chi_j$ that leads to the inequality $\mu_j^k \geqslant (\sum_\ell \pi_{j\ell}^k)\mu_j^k + \chi_j \Rightarrow (1 - \sum_\ell \pi_{j\ell}^k)\mu_j^k \geqslant \chi_j$. The proof follows from $(1 - \sum_\ell \pi_{j\ell}^k) > 0$ (see Definitions 4 and 6).

(ii) The proof is similar to that of the part (i). □

Proposition 4. *Given $\Pi(S^0) = \Pi^0 \equiv \Pi^{\text{plant}}$ and $\mu^k \equiv [I - \Pi^k]^{-1}X$, let Π^{k+1} be generated from Π^k for $k \geqslant 0$ as follows: $\forall i, j \in \{1, 2, \ldots, n\}$, ijth element of Π^{k+1} is modified as*

$$\pi_{ij}^{k+1} \begin{cases} \geqslant \pi_{ij}^k & \text{if } \mu_j^k > 0 \\ = \pi_{ij}^k & \text{if } \mu_j^k = 0 \\ \leqslant \pi_{ij}^k & \text{if } \mu_j^k < 0 \end{cases} \quad \text{and} \quad \Pi^k \leqslant \Pi^0 \quad \forall k.$$

Then, $\mu^{k+1} \geqslant \mu^k$ elementwise and equality holds if and only if $\Pi^{k+1} = \Pi^k$.

Proof.

$$\begin{aligned} \mu^{k+1} - \mu^k &= ([I - \Pi^{k+1}]^{-1} - [I - \Pi^k]^{-1})X \\ &= [I - \Pi^{k+1}]^{-1}([I - \Pi^k] - [I - \Pi^{k+1}])[I - \Pi^k]^{-1}X \\ &= [I - \Pi^{k+1}]^{-1}(\Pi^{k+1} - \Pi^k)\mu^k. \end{aligned}$$

Defining the matrix $\Delta^k \equiv \Pi^{k+1} - \Pi^k$, let the jth column of Δ^k be denoted as Δ_j^k. Then, $\Delta_j^k \leqslant 0$ if $\mu_j^k < 0$ and $\Delta_j^k \geqslant 0$ if $\mu_j^k > 0$, and the remaining columns of Δ^k are zero vectors. This implies that: $\Delta^k\mu^k = \sum_j \Delta_j^k \mu_j^k \geqslant 0$. Since $\Pi^k \leqslant \Pi^0\ \forall k$, $[I - \Pi^{k+1}]^{-1} \geqslant 0$ elementwise; we have $[I - \Pi^{k+1}]^{-1}\Delta^k\mu^k \geqslant 0 \Rightarrow \mu^{k+1} \geqslant \mu^k$. For $\mu_j^k \neq 0$ and Δ^k as defined above, $\Delta^k\mu^k = 0$ if and only if $\Delta^k = 0$. Then, $\Pi^{k+1} = \Pi^k$ and $\mu^{k+1} = \mu^k$. □

Corollary 1 to Proposition 4. Let $\mu_j^k < 0$. Let Π^{k+1} be generated from Π^k by disabling controllable events that lead to the state q_j. Then, $\mu_j^{k+1} < 0$.

Proof. Since only jth column of $[I - \Pi^{k+1}]$ is different from that of $[I - \Pi^k]$ and the remaining columns are the same, the jth row of the cofactor matrix of $[I - \Pi^{k+1}]$ is the same as that of the cofactor matrix of $[I - \Pi^k]$, we have $Det[I - \Pi^{k+1}]\mu_j^{k+1} = Det[I - \Pi^k]\mu_j^k$. By Corollary 2 to Proposition 2, both determinants are real positive. □

Remark 1. In Proposition 4, some elements of the jth column of Π^k are decreased (or increased) by disabling (or re-enabling) controllable events that lead to the states q_j for which $\mu_j^k < 0$ (or $\mu_j^k \geqslant 0$).

Proposition 5. *The iterations in Proposition 4 lead to a cost matrix Π^* that is optimal in the sense of maximizing the performance vector $\mu^* \equiv [I - \Pi^*]^{-1}\mathrm{X}$ elementwise.*

Proof. Let us consider an arbitrary cost matrix $\tilde{\Pi} \leqslant \Pi^0$ and $\tilde{\mu} \equiv [I - \tilde{\Pi}]^{-1}\mathrm{X}$. We will show that $\tilde{\mu} \leqslant \mu^*$. Let us rearrange the elements of the μ^*-vector such that $\mu^* = [\underbrace{\mu_1^* \cdots \mu_\ell^*}_{\geqslant 0} | \underbrace{\mu_{\ell+1}^* \cdots \mu_n^*}_{<0}]^{\mathrm{T}}$ and the cost matrices $\tilde{\Pi}$ and Π^* are also rearranged in the order in which the μ^*-vector is arranged.

According to Proposition 4, no controllable event leading to states q_k, $k = 1, 2, \ldots, \ell$ has been disabled; and all controllable events leading to states q_k, $k = \ell+1, \ell+2, \ldots, n$ have been disabled. Therefore, the elements in the first ℓ columns of Π^* are the same as those of the Π^0-matrix and only the elements in the last $(n-\ell)$ columns are decreased to the maximum permissible extent by disabling all controllable events. In contrast, the columns of $\tilde{\Pi}$ are reduced by an arbitrary choice. Since $\tilde{\mu} - \mu^* = [I - \tilde{\Pi}]^{-1}[\tilde{\Pi} - \Pi^*]\mu^*$, the $(\tilde{\Pi} - \Pi^*)$-matrix whose first ℓ columns are non-positive and last $(n-\ell)$ columns are non-negative yields:

$$\tilde{\mu} - \mu^* = [I - \tilde{\Pi}]^{-1}[\textit{first } \ell \textit{ cols} \leqslant 0 \,|\, \textit{last } (n-\ell) \times \textit{cols} \geqslant 0]\mu^*,$$

where

$$\mu^* = [\underbrace{\mu_1^* \cdots \mu_\ell^*}_{\geqslant 0} \,|\, \underbrace{\mu_{\ell+1}^* \cdots \mu_n^*}_{<0}]^{\mathrm{T}}.$$

Since $[I - \tilde{\Pi}]^{-1} \geqslant 0$ elementwise, we conclude that

$$\tilde{\mu} - \mu^* = \underbrace{[I - \tilde{\Pi}]^{-1}}_{\geqslant 0}\left(\underbrace{\sum_{j=1}^{\ell} Col_j \cdot \mu_j^*}_{\leqslant 0} + \underbrace{\sum_{j=\ell+1}^{n} Col_j \cdot \mu_j^*}_{\leqslant 0}\right) \leqslant 0.$$

Therefore, $\tilde{\mu} \leqslant \mu^*$ for any arbitrary choice of $\tilde{\Pi}$. □

Proposition 6. *The control policy induced by the Π^*-matrix is unique in the sense that the controlled language is most permissive (i.e., least restrictive) among all controller(s) having the best performance.*

Proof. Disabling controllable event(s) leading to a state q_j with performance measure $\mu_j^* = 0$ does not alter the performance vector μ^*. The optimal control does not disable any controllable event leading to the states with zero performance. This implies that, among all controllers with the identical performance μ^*, the control policy induced by the Π^*-matrix is most permissive. □

Remark 2. Propositions 5 and 6 suffice to conclude that the Π^*-matrix yields the most permissive controller with the best performance μ^*. The control policy is realized as:

- All controllable events leading to the states q_j, for which $\mu_j^* < 0$, are disabled;
- All controllable events leading to the states q_j, for which $\mu_j^* \geqslant 0$, are enabled.

4. Construction of the optimal control law

This section summarizes the construction of the optimal control law that maximizes the performance of the controlled language of the DFSA for any initial state $q \in Q$. Let G be a DFSA plant model without any constraint (i.e., operational specifications) and have the state transition cost matrix of the unsupervised plant as: $\Pi^{\mathrm{plant}} \in \mathfrak{R}^{n \times n}$ and the characteristic vector as: $\mathrm{X} \in \mathfrak{R}^n$. Then, the performance vector at $k = 0$ is given as $\mu^0 = [\mu_1^0 \mu_2^0 \cdots \mu_n^0]^{\mathrm{T}} = (I - \Pi^0)^{-1}\mathrm{X}$, where the jth element μ_j^0 of the vector μ^0 is the performance of the language, with the initial state q_j. Then, $\mu_j^0 < 0$ implies that, if the state q_j is reached, then the plant will yield bad performance. Intuitively, the control system should attempt to prevent the automaton from reaching q_j by disabling all controllable events that lead to this state. Therefore, the optimal control algorithm starts with disabling all controllable events that lead to every state q_j for which $\mu_j^0 < 0$. This is equivalent to reducing all elements of the corresponding columns of the Π^0-matrix by disabling the controllable events. In the next iteration, i.e., $k = 1$, the updated cost matrix Π^1 is obtained as: $\Pi^1 = \Pi^0 - \Delta^0$ where $\Delta^0 \geqslant 0$ (the inequality being implied elementwise) is composed of event costs corresponding to all controllable events that have been disabled. Using Proposition 4, $\mu^0 \leqslant \mu^1 \equiv [I - \Pi^1]^{-1}\mathrm{X}$. Although all controllable events leading to every state corresponding to a negative element of μ^1 are disabled, some of the controllable events that were disabled at $k = 0$ may now lead to states corresponding to positive elements of μ^1. Performance could be further enhanced by re-enabling these controllable events. For $k \geqslant 1$, $\Pi^{k+1} = \Pi^k + \Delta^k$ where $\Delta^k \geqslant 0$ is composed of all re-enabled controllable events at k.

If $\mu^0 \geqslant 0$, i.e., there is no state q_j such that $\mu_j^0 < 0$, then the plant performance cannot be improved by event disabling consequently, the null controller S^0 is the optimal controller for the given plant. Therefore, we consider the cases where $\mu_j^0 < 0$ for some state q_j.

Starting with $k = 0$ and $\Pi^0 \equiv \Pi^{\mathrm{plant}}$, the control policy is constructed by the following two-step procedure:

Step 1: For every state q_j for which $\mu_j^0 < 0$, disable controllable events leading to q_j. Now, $\Pi^1 = \Pi^0 - \Delta^0$, where $\Delta^0 \geqslant 0$ is composed of event costs corresponding to all controllable events, leading to q_j for which $\mu_j^0 < 0$, which have been disabled at $k = 0$.

Step 2: Starting with $k = 1$, if $\mu_j^k \geqslant 0$, re-enable all controllable events leading to q_j, which were disabled in Step 1. The cost matrix is updated as: $\Pi^{k+1} = \Pi^k + \Delta^k$ for $k \geqslant 1$, where $\Delta^k \geqslant 0$ is composed of event costs corresponding to all currently re-enabled controllable events. The iteration is terminated if no controllable event leading to q_j remains disabled for which $\mu_j^k > 0$. At this stage, the optimal performance $\mu^* \equiv [I - \Pi^*]^{-1}\mathrm{X}$.

Proposition 7. *The number of iterations needed to arrive at the optimal control law does not exceed the number*, n, *of states of the DFSA.*

Proof. Following Proposition 4, the sequence of performance vectors $\{\Pi^k\}$ in successive iterations of the two step procedure is monotonically increasing. The first iteration at $k = 0$ disables controllable events following Step 1 of the two-step procedure. During each subsequent iteration in Step 2, the controllable events leading to at least one state are re-enabled. When Step 2 is terminated, there remains at least one negative element, $\mu_j^k < 0$ by Corollary 1 to Proposition 4. Therefore, as the number of iterations in Step 2 is at most $n - 1$, the total number of iterations to complete the two-step procedure does not exceed n. □

Remark 3. Since each iteration requires a single Gaussian elimination of n unknowns from n linear algebraic equations, computational complexity of the control algorithm is polynomial in n.

5. An application example

As an example of state-based optimal control, this section presents the design of a discrete-event supervisor for a twin-engine unmanned aircraft that is used for surveillance and data collection. Engine health and operating conditions are monitored in real time based on observed data. In the event of any observed abnormality, the supervisor may decide to continue or abort the mission. Engine health and operating conditions, which are monitored in real time base on observed data, classified into three mutually exclusive and exhaustive categories: good; unhealthy (but operable); and inoperable.

The plant automaton model has 13 states (excluding the dump state), of which three are marked states, and nine events, of which four are controllable. All events are assumed to be observable. The states and events of the plant model are listed in Tables 1 and 2, respectively.

Based on the information provided in Tables 2, 3 and 4, an optimal control policy has been synthesized following the two-step procedure in Section 4. The values of the performance vectors, μ^0 for the unsupervised plant, μ^1 in Step 1, and μ^2 in Step 2, are summarized in Table 5. Clearly, $\mu^1 \geqslant \mu^0$ because controllable event(s), if any, leading to states 4 to 10 and 13 (for which $\mu_j^0 < 0$) are disabled in Step 1. As indicated by the dashed and dotted arcs in the state transition diagram of Fig. 1, the controllable event **k** leading to states 4, 7, 8, and 10 are disabled in Step 1 while the uncontrollable events were kept enabled as the supervisor has no authority to disable them. The event **k** leading to the state 4 is re-enabled in Step 2.

Table 1
Plant automaton states

State#	State description
1	Safe in base
2	Mission executing—two good engines
3	One engine unhealthy
4	Mission executing—one good and one unhealthy engine
5	Both engines unhealthy
6	One engine good and one engine inoperative
7	Mission execution with two unhealthy engines
8	Mission execution with only one good engine
9	One engine unhealthy and one engine inoperative
10	Mission execution with only one unhealthy engine
11	Mission aborted or not completed (*Bad Marked State*)
12	Mission successful (*Good Marked State*)
13	Aircraft destroyed (*Bad Marked State*)

Table 2
Plant event alphabet

Event#	Event description
s	start and take-off (*Controllable*)
b	one good engine becoming unhealthy
t	one unhealthy engine becoming inoperative
v	one good engine becoming inoperative
k	keep engine(s) running (*Controllable*)
a	mission abortion (*Controllable*)
f	mission completion
d	destroyed aircraft
ℓ	landing (*Controllable*)

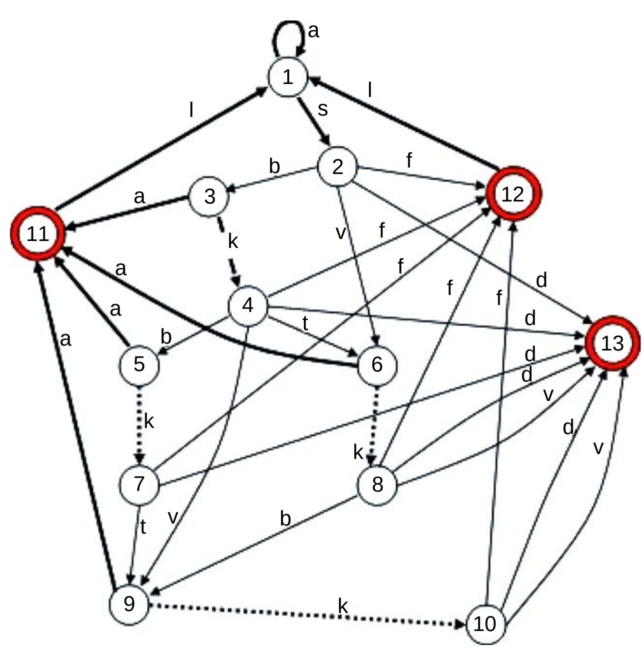

Fig. 1. Plant state transitions.

Table 3
State transition and event cost matrix

	s	b	t	v	k	a	f	d	ℓ
1	0.5 (2)					0.02 (1)			
2		0.05 (3)		0.01 (6)			0.8 (12)	0.1 (13)	
3					0.45 (4)	0.45 (11)			
4		0.12 (5)	0.16 (6)	0.1 (9)			0.5 (12)	0.12 (13)	
5					0.45 (7)	0.45 (11)			
6					0.45 (8)	0.45 (11)			
7			0.25 (9)				0.5 (12)	0.2 (13)	
8		0.2 (9)		0.01 (13)			0.3 (12)	0.4 (13)	
9					0.45 (10)	0.45 (11)			
10			0.35 (13)				0.2 (12)	0.40 (13)	
11									0.95 (1)
12									0.95 (1)
13									

Table 4
State transition cost matrix Π^{plant}

0.02	0.50	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.05	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.80	0.10
0.00	0.00	0.00	0.45	0.00	0.00	0.00	0.00	0.00	0.00	0.45	0.00	0.00
0.00	0.00	0.00	0.00	0.12	0.16	0.00	0.00	0.10	0.00	0.00	0.50	0.12
0.00	0.00	0.00	0.00	0.00	0.00	0.45	0.00	0.00	0.00	0.45	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.45	0.00	0.00	0.45	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.25	0.00	0.00	0.50	0.20
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.20	0.00	0.00	0.30	0.41
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.45	0.45	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.20	0.75
0.95	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.95	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Characteristic Vector $X = [0\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ -0.05\ \ 0.25\ -1.0]^{\mathrm{T}}$

Table 5
Performance (μ) vector

Initial state	μ^0 (Unsupervised)	μ^1 (Step 1)	$\mu^* = \mu^2$ (Step 2)
1	0.0823	0.0840	0.0850
2	0.1613	0.1646	0.1665
3	0.0062	0.0134	0.0366
4	**−0.0145**	**0.0500**	0.0506
5	−0.0367	0.0134	0.0138
6	−0.1541	0.0134	0.0138
7	−0.1097	−0.0317	−0.0312
8	−0.3706	−0.3084	−0.3080
9	−0.2953	0.0134	0.0138
10	−0.6844	−0.6840	−0.6839
11	0.0282	0.0298	0.0307
12	0.3282	0.3298	0.3307
13	−1.0000	−1.0000	−1.0000

The state transition function δ and the entries $\tilde{\pi}_{ij}$ (see Definition 3) are entered simultaneously in Table 3. The fraction part in each entry denotes the corresponding state-based event cost $\tilde{\pi}_{ij}$ such that each row sum of the event cost matrix $\tilde{\Pi}$ is strictly less than one. The integer part (within parentheses) in each entry denotes the respective destination resulting from the occurrence of the event. The values of $\tilde{\pi}_{ij}$ were selected by extensive simulation experiments on gas turbine engine models and were also based on experience of gas turbine engine operation and maintenance. The dump state and any transitions to the dumped state are not shown in Table 3. The elements of the characteristic vector (see Definition 2) are chosen as weights ranging between −1 and 1 based on the perception of each marked state's role on the gas turbine system performance. The information provided in Table 3 is sufficient to generate the Π^{plant}-matrix that is listed in Table 4 which also includes the characteristic vector X.

Note that $\mu_4^1 > 0$ whereas $\mu_4^0 < 0$ in Table 5. As indicated by the dashed arc in Fig. 1, the controllable event **k** leading to state 4 (which was disabled in Step 1) is now re-enabled in Step 2. Consequently, the performance vector is further increased, i.e., $\mu^2 \geqslant \mu^1$. Since there is no sign change between the elements of μ^1 and μ^2, the procedure is terminated after $k = 2$ following Step 2 and the resulting optimal performance vector is $\mu^* = \mu^2$. Table 5 shows that if the initial state is 1, then the best achievable performance of any controlled plant is 0.0850 based on the language measure parameters in Table 4. The controller also yields the most permissive controllable sublanguage that achieves the best performance $\mu^* = \mu^2$.

The best performance μ^* of the above optimal controller is also verified by comparison with three other controllers that were designed independently using the following specifications:

- Specification# 1: at least one of the two engines must be in good condition for mission continuation.

- Specification# 2: both engines must be in operable condition for mission continuation.
- Specification# 3: both engines must be in good condition for mission continuation.

The above three controllers were designed using a Java-based graphical interactive package that is coded following the supervisory control theory of Ramadge and Wonham (1987). The performance measure μ_1 (i.e., with the initial state 1) of the unsupervised plant is 0.0823, and that of three controllers with specification# 1, 2, and 3 is evaluated to be: 0.0826, 0.0843, and 0.0840, respectively. The performance of each of these controllers is inferior to the performance, 0.0850, of the optimal controller as expected.

6. Summary and conclusions

This paper presents the theory, formulation, and validation of an unconstrained optimal control policy for finite state automata that may have already been subjected to constraints such as control specifications. The synthesis procedure is quantitative and relies on a signed real measure of formal languages, which is based on a specified state transition cost matrix and a characteristic vector (Wang & Ray, 2002). The state-based optimal control policy maximizes the performance vector by selectively disabling controllable events that may lead to bad marked states and simultaneously ensuring that the remaining controllable events are kept enabled. The goal is to maximize the measure of the controlled plant language without any further constraints. The control policy induced by the updated state transition cost matrix yields maximal performance and is unique in the sense that the controlled language is most permissive (i.e., least restrictive) among all controller(s) having the optimal performance.

Derivation of the control policy requires at most n iterations, where n is the number of states of the DFSA model and each iteration is required to solve a set of n simultaneous linear algebraic equations. As such computational complexity of the control synthesis procedure is polynomial in the number of states. The procedure for synthesis of the optimal control algorithm has been validated on a finite-state machine model of gas turbine engine operations.

Future areas of research in optimal control include: (i) incorporation of the cost of disabling controllable events (Fu, Ray, & Lagoa, 2003); and (ii) robustness of the control policy relative to unstructured and structured uncertainties in the plant model including variations in the language measure parameters.

References

Brave, Y., & Heyman, M. (1990). *On optimal attraction of discrete-event processes*. Technical Report, Technion, Haifa, October 1990.

Fu, J., Ray, A., & Lagoa, C. (2003). Optimal control of regular languages with event disabling cost. Preprints; *Proceedings of the American control conference*, Denver, Colorado, June 2003.

Kumar, R., & Garg, V. (1995). Optimal supervisory control of discrete event dynamical systems. *SIAM Journal on Control and Optimization*, *33*(2), 419–439.

Naylor, A. W., & Sell, G. R. (1982). *Linear operator theory in science and engineering*. New York: Springer.

Pasino, K. M., & Antsaklis, P. J. (1989). On the optimal control of discrete event systems. *IEEE decision and control conference*, Tampa, Florida (pp. 2713–2718).

Ramadge, P. J., & Wonham, W. M. (1987). Supervisory control of a class of discrete event processes. *SIAM Journal on Control and Optimization*, *25*(1), 206–230.

Sengupta, R., & Lafortune, S. (1998). An optimal control theory for discrete event systems. *SIAM Journal on Control and Optimization*, *36*(2), 488–541.

Surana, A., & Ray, A. (2003). Signed real measure of regular languages. *42nd IEEE conference on decision and control* (*CDC*), Maui, Hawaii, December 2003.

Wang, X., & Ray, A. (2002). Signed real measure of regular languages. *American control conference*, Anchorage, Alaska, May 2002 (pp. 3937–3942).

Jinbo Fu received the Bachelor and Master degrees, both in Precision Instruments and Mechanology, from Tsinghua University, Beijing, China, in 1996 and 1999, respectively. He also received Master degrees in Electrical Engineering and Mechanical Engineering from Pennsylvania State University, University Park in 2002 and 2003 respectively, and the Ph.D. degree in Mechanical Engineering from Pennsylvania State University, University Park in 2003.

He has been working on the theories and applications of discrete event supervisory (DES) control since 2000 in several projects. His research interests include DES control theories and application, intelligent systems, engine propulsion system command and control, damage mitigating control, fault detection and information fusion.

Asok Ray earned the Ph.D. degree in Mechanical Engineering from Northeastern University, Boston, MA, and also graduate degrees in each discipline of Electrical Engineering, Mathematics, and Computer Science. Dr. Ray joined the Pennsylvania State University in July 1985, and is currently a Distinguished Professor of Mechanical Engineering. Prior to joining Penn State, Dr. Ray held research and academic positions at Massachusetts Institute of Technology and Carnegie-Mellon University as well as research and management positions at GTE Strategic Systems Division, Charles Stark Draper Laboratory, and MITRE Corporation. Dr. Ray has been a Senior Research Fellow at NASA Glenn Research Center under a National Academy of Sciences award. Dr. Ray's research experience and interests include: Control and optimization of continuously varying and discrete-event dynamical systems; Intelligent instrumentation for real-time distributed systems; and Modeling and analysis of complex dynamical systems from thermodynamic perspectives in both deterministic and stochastic settings, as applied to Aeronautics and Astronautics, Undersea Vehicles and Surface Ships, Power and Processing plants, and Robotics. Dr. Ray has authored or co-authored three hundred seventy research publications including over one hundred and fifty scholarly articles in refereed journals such as transactions of ASME, IEEE and AIAA, and research monographs. Dr. Ray is a Fellow of IEEE, a Fellow of ASME, and an Associate Fellow of AIAA.

Constantino M. Lagoa got his B.S. and M.Sc. degrees from the Instituto Superior Tecnico, Technical University of Lisbon, Portugal in 1991 and 1994, respectively, and his Ph.D. degree from the University of Wisconsin at Madison in 1998. He joined the Electrical Engineering Department of the Pennsylvania State University in August 1998, where he currently holds the position of Assistant Professor. He has a wide range of research interests including robust control, controller design under risk specifications, control of computer networks and discrete event dynamical systems. In 2000 he received the NSF CAREER award for his proposal on system design under risk constraints.